



The Gravity Model

James E. Anderson

Department of Economics, Boston College, Chestnut Hill, Massachusetts 02467, and
NBER, Cambridge, Massachusetts 02138; email: james.anderson.1@bc.edu

Annu. Rev. Econ. 2011. 3:133–60

The *Annual Review of Economics* is online at
economics.annualreviews.org

This article's doi:
[10.1146/annurev-economics-111809-125114](https://doi.org/10.1146/annurev-economics-111809-125114)

Copyright © 2011 by Annual Reviews.
All rights reserved

JEL codes: F10, F22, J61, F10, R40

1941-1383/11/0904-0133\$20.00

Keywords

incidence, multilateral resistance, trade costs, migration

Abstract

Gravity has long been one of the most successful empirical models in economics. Incorporating deeper theoretical foundations of gravity into recent practice has led to a richer and more accurate estimation and interpretation of the spatial relations described by gravity. Wider acceptance has followed. Recent developments are reviewed here, and suggestions are made for promising future research.

1. INTRODUCTION

The gravity model in economics was until relatively recently an intellectual orphan, unconnected to the rich family of economic theory. This review is a tale of the orphan's reunion with its heritage and the benefits that continue to flow from connections to more distant relatives.

Gravity has long been one of the most successful empirical models in economics, ordering remarkably well the enormous observed variation in economic interaction across space in both trade and factor movements. The good fit and relatively tight clustering of coefficient estimates in the vast empirical literature suggested that some underlying economic law must be at work, but in the absence of an accepted connection to economic theory, most economists ignored gravity. The authoritative survey of Leamer & Levinsohn (1995) captures the state of professional thinking in the mid-1990s: "These estimates of gravity have been both singularly successful and singularly unsuccessful. They have produced some of the clearest and most robust empirical findings in economics. But, paradoxically, they have had virtually no effect on the subject of international economics. Textbooks continue to be written and courses designed without any explicit references to distance, but with the very strange implicit assumption that countries are both infinitely far apart and infinitely close, the former referring to factors and the latter to commodities." Subsequently, gravity first appeared in textbooks in 2004 (Feenstra 2004), following success in connecting gravity to economic theory, the subject of Section 4.

Reviews are not intended to be surveys. My take on the gravity model, thus licensed to be idiosyncratic, scants or omits some topics that others have found important while it emphasizes some topics that others have scanted. My emphases and omissions are intended to guide the orphan to maturity. An adoptive parent's biases may have contaminated my judgment, *caveat emptor*.

My focus is on theory. Incorporating the theoretical foundations of gravity into recent practice has led to richer and more accurate estimation and interpretation of the spatial relations described by gravity, so where appropriate I point out this benefit. The harvest reaped from empirical work applying the gravity model is recently surveyed elsewhere (Anderson & van Wincoop 2004, Bergstrand & Egger 2011).

From a modeling standpoint, gravity is distinguished by its parsimonious and tractable representation of economic interaction in a many-country world. Most international economic theory is concentrated on two-country cases, occasionally extended to three-country cases with special features. The tractability of gravity in the many-country case results from its modularity: The distribution of goods or factors across space is determined by gravity forces conditional on the size of economic activities at each location. Modularity readily allows for disaggregation by goods or regions at any scale and permits inference about trade costs that is not dependent on any particular model of production and market structure in full general equilibrium. The modularity theme recurs often below but is missing from some other prominent treatments of gravity in the literature.

2. TRADITIONAL GRAVITY

The story begins by setting out the traditional gravity model and noting clues to its union with economic theory. The traditional gravity model drew on analogy with

Newton's law of gravitation. A mass of goods or labor or other factors of production supplied at origin i , Y_i , is attracted to a mass of demand for goods or labor at destination j , E_j , but the potential flow is reduced by the distance between them, d_{ij} . Strictly applying the analogy,

$$X_{ij} = Y_i E_j / d_{ij}^2$$

gives the predicted movement of goods or labor between i and j , X_{ij} . Ravenstein (1889) pioneered the use of gravity for migration patterns in the nineteenth-century United Kingdom. Tinbergen (1962) was the first to use gravity to explain trade flows. Departing from strict analogy, traditional gravity allowed the exponents of 1 applied to the mass variables and of 2 applied to bilateral distance to be generated by data to fit a statistically inferred loglinear relationship between data on flows and the mass variables and distance.

Generally, across many applications, the estimated coefficients on the mass variables cluster close to 1 and the distance coefficients cluster close to -1 , while the estimated equation fits the data well: Most data points cluster close to the fitted line in the sense that 80%–90% of the variation in the flows is captured by the fitted relationship. The fit of traditional gravity improved when supplemented with other proxies for trade frictions, such as the effect of political borders and common language.

Notice that bilateral frictions alone would appear to be inadequate to fully explain the effects of trade frictions on bilateral trade because the sale from i to j is influenced by the resistance to movement on i 's other alternative destinations and by the resistance on movement to j from j 's alternative sources of supply. Prodded by this intuition, the traditional gravity literature recently developed remoteness indexes of each country's average effective distance to or from its partners ($\sum_i d_{ij}/Y_i$ was commonly defined as the remoteness of country j) and used them as further explanatory variables in the traditional gravity model, with some statistical success.

The general problem posed by the intuition behind remoteness indexes is analogous to the N-body problem in Newtonian gravitation. An economic theory of gravity is required for an adequate solution. Because there are many origins and many destinations in any application, a theory of the bilateral flows must account for the relative attractiveness of origin-destination pairs. Each sale has multiple possible destinations, and each purchase has multiple possible origins: Any bilateral sale interacts with all others and involves all other bilateral frictions. This general equilibrium problem is neatly solved with structural gravity models. For expositional ease, the discussion focuses below on goods movements, except when migration or investment is specifically treated.

3. FRICTIONLESS GRAVITY LESSONS

Taking a step toward structure, an intuitively appealing starting point is the description of a completely smooth homogeneous world in which all frictions disappear. Developing the implications of this structure yields a number of useful insights about the pattern of world trade.

A frictionless world implies that each good has the same price everywhere. In a homogeneous world, economic agents everywhere might be predicted to purchase goods in the same proportions when faced with the same prices. In the next section, the assumptions on

preferences and/or technology that justify this plausible prediction are the focus, but here the focus is on the implications for trade patterns. In a completely frictionless and homogeneous world, the natural benchmark prediction is that $X_{ij}/E_j = Y_i/Y$, where the proportion of spending by j on goods from i is equal to the global proportion of spending on goods from i , and Y denotes world spending.

Any theory must impose adding up constraints, which for goods requires that the sum of sales to all destinations must equal Y_i , the total sales by origin i , and the sum of purchases from all origins must equal E_j , the total expenditure for each destination j . Total sales and expenditures must be equal: i.e., $\sum_i Y_i = \sum_j E_j = Y$.

One immediate payoff is an implication for inferring trade frictions. Multiplying both sides of the frictionless benchmark prediction $X_{ij}/E_j = Y_i/Y$ by E_j yields predicted frictionless trade $Y_i E_j / Y$. The ratio of observed trade X_{ij} to predicted frictionless trade $Y_i E_j / Y$ represents the effect of frictions along with random influences. (Bilateral trade data are notoriously rife with measurement error.) Fitting the statistical relationship between the ratio of observed to frictionless trade and various proxies for trade costs is justified by this simple theoretical structure as a proper focus of empirical gravity models.

Thus far, the treatment of trade flows has been of a generic good that most of the literature has implemented as an aggregate: the value of aggregate bilateral trade in goods, for example. But the model applies more naturally to disaggregated goods (and factors) because the frictions to be analyzed below are likely to differ markedly by product characteristics. The extension to disaggregated goods, indexed by k , is straightforward:

$$X_{ij}^k = \frac{Y_i^k E_j^k}{Y^k} = s_i^k b_j^k Y^k. \quad (1)$$

Here $s_i^k = Y_i^k / Y^k$ is country i 's share of the world's sales of goods class k , and $b_j^k = E_j^k / Y^k$ is country j 's share of the world's spending on k , equal to Y^k , the world's sales of k .

The notation and logic also readily apply to the disaggregation of countries into regions, and indeed a prominent portion of the empirical literature has examined bilateral flows between city pairs or regions, motivated by the observation that much economic interaction is concentrated at very short distances. The model can be interpreted to reflect individual decisions aggregated with a probability model (see Section 6.1 below).

In aggregate gravity applications (i.e., most applications), it has been common to use origin and destination mass variables equal to gross domestic product (GDP). This is conceptually inappropriate and leads to inaccurate modeling unless the ratio of gross shipments to GDP is constant (in which case the ratio goes into a constant term). A possible direction for aggregate modeling is to convert trade to the same value-added basis as GDP, but this seems more problematic than using disaggregated gravity to explain the pattern of gross shipments and then uniting estimated gravity models within a superstructure to connect to GDP. That is the strategy of the structural gravity model research program reviewed here.

Equation 1 generates a number of useful implications. (a) Big producers have big market shares everywhere. (b) Small sellers are more open in the sense of trading more with the rest of the world. (c) The world is more open the more similar in size, and the more specialized, the countries are. (d) The world is more open the greater the number of

countries. (e) World openness rises with convergence under the simplifying assumption of balanced trade.

Implication (a), that big producers have big market shares everywhere, follows because, reverting to the generic notation and omitting the k superscript, the frictionless gravity prediction is that

$$X_{ij}/E_j = s_i.$$

Implication (b), that small sellers are more open in the sense of trading more with the rest of the world, follows from

$$\sum_{i \neq j} X_{ij}/E_j = 1 - Y_j/Y = 1 - s_j$$

using $\sum_j E_j = \sum_i Y_i$, which is balanced trade for the world.

Implication (c) is that the world is more open the more similar in size, and the more specialized, the countries are. It is convenient to define world openness as the ratio of international shipments to total shipments, $\sum_j \sum_{i \neq j} X_{ij}/Y$. Dividing Equation 1 through by Y^k and suppressing the goods index k , world openness is given by

$$\sum_j \sum_{i \neq j} X_{ij}/Y = \sum_j b_j(1 - s_j) = 1 - \sum_j b_j s_j.$$

Using standard statistical properties,

$$\sum_j b_j s_j = Nr_{bs} \sqrt{\text{Var}(s)\text{Var}(b)} + 1/N,$$

where N is the number of countries or regions, Var denotes variance, r_{bs} is the correlation coefficient between b and s , and $1/N = \sum_i s_i/N = \sum_j b_j/N$, the average share. This equation follows from the shares summing to one and using standard properties of covariance. Here, $\text{Var}(s)$ and $\text{Var}(b)$ measure size dissimilarity, and the correlation of s and b , r_{bs} , is an inverse measure of specialization. Substituting into the expression for world openness,

$$\sum_j \sum_{i \neq j} X_{ij}/Y = 1 - 1/N - Nr_{bs} \sqrt{\text{Var}(s)\text{Var}(b)}. \quad (2)$$

Implication (c) follows from Equation 2 because on the right-hand side the similarity of country size shrinks the variances while specialization shrinks the correlation r_{bs} .

The country-size similarity property has been prominently stressed in the monopolistic competition and trade literature. (It is sometimes taken as evidence for monopolistic competition in a sector rather than as a consequence of gravity no matter what explains the pattern of the b 's and s 's.) The specialization property has also been noted in that literature as reflecting forces that make for greater net international trade, the absolute value of $s_j - b_j$. Making comparisons across goods classes, variation in the right-hand side of Equation 2 results from variation in specialization and in the dispersion of the shipment and expenditure shares. Notice again that the cross-commodity variation in world openness arises here in a frictionless world, a reminder that measures of world home bias in a world with frictions must be evaluated relative to the frictionless world benchmark.

Country-size similarity also tends to increase bilateral trade between any pair of countries, all else equal. This point (Bergstrand & Egger 2007) is seen most clearly with aggregate trade that is also balanced, hence $s_j = b_j$. Equation 1 can be rewritten as

$$X_{ij} = s_i^{ij} s_j^{ij} \frac{(Y_i + Y_j)^2}{Y},$$

where $s_i^{ij} \equiv Y_i/(Y_i + Y_j)$, the share of i in the joint GDP of i and j . The product $s_i^{ij} s_j^{ij}$ is maximized at $s_i^{ij} = s_j^{ij} = 1/2$, so for a given joint GDP size, bilateral trade is increasing in country similarity. (With unbalanced trade or specialization, an analogous similarity property holds for the bilateral similarity of income and expenditure shares. Let $\gamma_j = E_j/Y_j$. Then the same equation as before holds with the right-hand side multiplied by γ_j .)

A more novel implication of Equation 2 is implication (d), that world openness is ordinarily increasing in the number of countries. Increasing world openness due to a rise in the number of countries reflects the property that smaller countries are more naturally open and division makes for more and smaller countries.

This effect is seen by differentiating the left-hand side of $\sum_j \sum_{i \neq j} X_{ij}/Y = 1 - \sum_j b_j s_j$, yielding $-\sum_j (b_j ds_j + s_j db_j)$. Increasing the number of countries tends to imply reducing the share of each existing country while increasing the share (from zero) of the new country. The preceding differential expression should thus ordinarily be positive.

The qualification with the term ordinarily is needed because the pattern of share changes will depend on the underlying structure as revealed by the left-hand side of Equation 2. On the one hand, the average share $1/N$ decreases as N rises, raising world openness. On the other hand, the change in the number of countries will usually change $r_{bs} \sqrt{\text{Var}(b)\text{Var}(s)}$ in ways that depend on the type of country division (or confederation) as well as on indirect effects on shares as prices change. [The apparent direct effect of N in the first term on the right-hand side of Equation 2 vanishes because $1/N$ scales $\sqrt{\text{Var}(b)\text{Var}(s)}$.]

A practical implication of this discussion is that intertemporal comparisons of ratios of world international trade to world income, to be economically meaningful, should be controlled for changes in the size distribution and the number of countries, a correction of large practical importance in the past 50–100 years. Alternatively, measures of openness meant to reflect the effects of trade frictions should be constructed in relation to the frictionless benchmark.

Applied to aggregate trade data, gravity yields implication (e), that world openness rises with convergence under the simplifying assumption of balanced trade for each country, $b_j = s_j, \forall j$. The right-hand side of Equation 2 becomes $N\text{Var}(s) + 1/N$ under balanced trade, and per-capita income convergence lowers $\text{Var}(s)$ toward the variance of population. Baier & Bergstrand (2001) use the convergence property to partially explain postwar growth in world trade/income, finding relatively little action, although presumably more recent data influenced by the rise of China and India might give more action.

Pointing toward a connection with economic theory, the shares s_i and b_j and the plausible hypothesis of the frictionless model must originate from an underlying structure of preferences and technology. Also, the deviation of observed X_{ij} from the frictionless prediction reflects frictions as they act on the pattern of purchase decisions of buyers and the sales decisions of sellers, which originate from an underlying structure of preferences and technology.

4. STRUCTURAL GRAVITY

Modeling economies with trade costs works best if it moves backward from the end user. Start by evaluating all goods at user prices, applying demand-side structure to determine the allocation of demand at those prices. Treat all costs incurred between production and end use as being incurred by the supply side of the market, even though there are often significant costs directly paid by the user. What matters economically in the end is the full cost between production and end use, and the incidence of that cost on the producer and end user. Many of these costs are not directly observable, and the empirical gravity literature indicates the total is well in excess of the transportation and insurance costs that are observable (see Anderson & van Wincoop 2004 for a survey of trade costs).

The supply side of the market under this approach both produces and distributes the delivered goods, incurring resource costs that are paid by end users. The factor markets for those resources must clear at equilibrium factor prices, determining costs that link to end-user prices. Budget constraints require national factor incomes to pay for national expenditures plus net lending or transfers including remittances. Below the national accounts, individual economic agents also meet budget constraints. Goods markets clear when prices are found such that demand is equal to supply for each good. The full general equilibrium requires a set of bilateral factor prices and bilateral goods prices such that all markets clear and all budget constraints are met.

This standard description of general economic equilibrium is too complex to yield something like gravity. A hugely useful simplification is modularity, subordinating the economic determination of equilibrium distribution of goods within a class under the superstructure determination of the distribution of production and expenditure between classes of goods. Anderson & van Wincoop (2004) call this property trade separability. Observing that goods are typically supplied from multiple locations, even within fine census commodity classes, it is natural to look for a theoretical structure that justifies grouping in this way. The structural gravity model literature has uncovered two structures that work, one on the demand side and one on the supply side, detailed in Sections 4.1 and 4.2.

Modularity (trade separability) permits the analyst to focus exclusively on inference about distribution costs from the pattern of distribution of goods (or factors) without having to explain at the same time what determines the total supplies of goods to all destinations or the total demand for goods from all origins. This is a great advantage for two reasons. First, it simplifies the inference task enormously. Second, the inferences about the distribution of goods or factors are consistent with a great many plausible general equilibrium models of national (or regional) production and consumption.

Modularity also requires a restriction on trade costs, so that only the national aggregate burden of trade costs within a goods class matters for allocation between classes. The most popular way to meet this requirement is to restrict the trade costs so that the distribution of goods uses resources in the same proportion as the production of those same goods. Samuelson (1952) invented iceberg melting trade costs in which the trade costs were proportional to the volume shipped, as the amount melted from the iceberg is proportional to its volume. The iceberg metaphor still applies when allowing for a fixed cost, as if a chunk of the iceberg breaks off as it parts from the mother glacier. Mathematically, the generalized iceberg trade cost is linear in the volume shipped. Economically, distribution continues to require resources to be used in the same proportion as in production. Fixed

costs are realistic and potentially play an important role in explaining why many potential bilateral flows are equal to zero.

More general nonlinear trade cost functions continue to satisfy the production proportionality restriction and thus meet the requirements of modularity, but depart from the iceberg metaphor. Bergstrand (1985) derives a joint cost function that is homogeneous of degree one with constant elasticity of transformation. This setup allows for substitution effects in costs between destinations rather than the cost independence because of fixed coefficients in the iceberg model. Bilateral costs have a natural aggregator that is an iceberg cost facing monopolistically competitive firms. A nice feature of the joint cost model is its econometric tractability under the hypothesis of the profit-maximizing choice of destinations. Although potentially more realistic, the joint cost refinement turns out to make relatively little difference empirically.

Arkolakis (2008) develops a nonlinear (in volume) trade cost function in which heterogeneous customers are obtained by firms with a marketing technology featuring a fixed-cost component (running a national advertisement) and a variable-cost component (leafletting or telemarketing) subject to diminishing returns as the less likely customers are encountered. Because of the Ricardian production and distribution technology, resource requirements in distribution remain proportional to production resource requirements. Arkolakis shows that the marketing technology model can rationalize features of the firm-level bilateral shipments data that cannot be explained with the linear fixed-costs model. His setup is not econometrically tractable but is readily applicable as a simulation model.

In all applications based on the preceding cost functions, proxies for costs are entered in some convenient functional form, usually loglinear in variables such as bilateral distance, contiguity, membership of a country, continent or regional trade agreement, common language, and common legal traditions (see Anderson & van Wincoop 2004 for more discussion). More generality in trade costs that violates the production proportionality restriction comes at the price of losing modularity (see Matsuyama 2007 for a recent exploration of the implications of noniceberg trade costs in a two-country Ricardian model; see Deardorff 1980 for a general treatment of the resource requirements of trade costs as a setting for his demonstration that the law of comparative advantage holds quite generally).

4.1. Demand-Side Structure

The second requirement for modularity can be met by restricting the preferences and/or technology such that the cross effects in demand between classes of goods (either intermediate or final) flow only through aggregate price indexes. This demand property is satisfied when preferences or technologies are homothetic and weakly separable with respect to a partition into classes with members defined by location, a partition structure called the Armington assumption. Thus, for example, steel products from all countries are members of the steel class. Notice that the assumption implies that goods are purchased from multiple sources because they are evaluated differently by end users, and goods are differentiated by place of origin.

It is usual to impose identical preferences across countries. Differences in demand across countries, such as a home bias in favor of locally produced goods, can be accommodated, understanding that “trade costs” now include the effect of a demand-side home bias. In practice it is difficult to distinguish demand-side home bias from the effect of trade costs, as

the proxies used in the literature (e.g., common language, former colonial ties, or internal trade dummies) plausibly pick up both demand and cost differences. Henceforth, the term trade cost is used without qualification but is understood to potentially reflect demand-side home bias. Declines in trade costs can be understood as reflecting homogenization of tastes.

Separability implies that each goods class has a natural quantity aggregate and a natural price aggregate, with substitution between goods classes occurring as if the quantity aggregates were goods in the standard treatment. The separability assumption implies that national origin expenditure shares within the steel class are not altered by changes in the prices of nonsteel products, although, of course, the aggregate purchase of steel is affected by the aggregate cross effect. Homotheticity ensures that relative demands are functions only of relative aggregate prices.

The first economic foundation for the gravity model was based on specifying the expenditure function to be a constant elasticity of substitution (CES) function (Anderson 1979). Expenditure shares in the CES case are given by

$$\frac{X_{ij}}{E_j} = \left(\frac{\beta_i p_i t_{ij}}{P_j} \right)^{1-\sigma}, \quad (3)$$

where P_j is the CES price index, σ is the elasticity of substitution parameter, β_i is the distribution parameter for varieties shipped from i , p_i is their factory gate price, and $t_{ij} > 1$ is the trade cost factor between origin i and destination j . The CES price index is given by

$$P_j = \left(\sum_i (\beta_i p_i t_{ij})^{1-\sigma} \right)^{1/(1-\sigma)}. \quad (4)$$

Notice that the same parameters characterize expenditure behavior in all locations; preferences are common across the world by assumption. Notice also that the shares are invariant to income; preferences are homothetic. With frictionless trade, $t_{ij} = 1, \forall(i, j)$, and therefore all the buyers' shares of good i must equal the sellers' share of world sales (at destination prices), Y_i/Y . Thus the frictionless benchmark is justified by assuming identical homothetic preferences. For intermediate goods, the same logic works by replacing expenditure shares with cost shares.

The distribution parameters β_i bear several interpretations. They could be exogenous taste parameters. Alternatively, in applications to monopolistically competitive products, β_i is proportional to the number of firms from i offering distinct varieties (Bergstrand 1989). Countries with more active firms get bigger weights. In long-run monopolistic competition, the number of firms is endogenous. Owing to fixed entry costs, bigger countries have more active firms in equilibrium, all else equal. The number of active firms contributes to determining the Y_i 's that are given in the gravity module.

The other building block in the structural gravity model is market clearance: at delivered prices $Y_i = \sum_j X_{ij}$. Multiplying both sides of Equation 3 by E_j and summing over j yields a solution for $\beta_i p_i^{1-\sigma}$:

$$\beta_i p_i^{1-\sigma} = \frac{Y_i}{\sum_j (t_{ij}/P_j)^{1-\sigma} E_j}.$$

Define the denominator as $\Pi_i^{1-\sigma}$.

Substituting into Equations 3 and 4 yields the structural gravity model:

$$X_{ij} = \frac{E_j Y_i}{Y} \left(\frac{t_{ij}}{P_j \Pi_i} \right)^{1-\sigma}, \quad (5)$$

$$(\Pi_i)^{1-\sigma} = \sum_j \left(\frac{t_{ij}}{P_j} \right)^{1-\sigma} \frac{E_j}{Y}, \quad (6)$$

$$(P_j)^{1-\sigma} = \sum_i \left(\frac{t_{ij}}{\Pi_i} \right)^{1-\sigma} \frac{Y_i}{Y}. \quad (7)$$

The second ratio on the right-hand side of Equation 5 is a decreasing function (under the empirically valid restriction $\sigma > 1$) of direct bilateral trade costs relative to the product of two indexes of all bilateral trade costs in the system.

Anderson & van Wincoop (2003) call the terms P_j and Π_i inward and outward multilateral resistance, respectively. Note that $\{P_j^{1-\sigma}, \Pi_i^{1-\sigma}\}$ can be solved from Equations 6 and 7 for given $t_{ij}^{1-\sigma}$'s, E_j 's, and Y_i 's combined with a normalization.¹ Under the assumption of bilateral trade cost symmetry $t_{ij} = t_{ji}, \forall i, j$ and balanced trade $E_j = Y_j, \forall j$, the natural normalization is $\Pi_i = P_i$. Anderson & van Wincoop estimated their gravity equation for Canada's provinces and U.S. states with a full information estimator that utilized Equation 7 with $\Pi_i = P_i$. Subsequent research has focused mostly on estimating Equation 5 with directional country fixed effects to control for $E_j/P_j^{1-\sigma}$ and $Y_i/\Pi_i^{1-\sigma}$.

Multilateral resistance is, on the face of it, an index of inward and outward bilateral trade costs, but because of the simultaneity of the system (Equations 6 and 7), all bilateral trade costs in the world contribute to the solution values. This somewhat mysterious structure has a simple and intuitive interpretation: Inward and outward multilateral resistance measure average buyers' and sellers' incidence of trade costs, respectively.

The incidence interpretation follows because the uniform preferences assumption in demand implies that the seller in effect makes a single shipment at a uniform markup factor Π_i to a world market with a share determined by

$$\frac{Y_i}{Y} = \left(\frac{\beta_i p_i \Pi_i}{P_W} \right)^{1-\sigma}. \quad (8)$$

The right-hand side of Equation 8, referring to the general form given in Equation 3, is interpreted as the global expenditure share on the good from i in a hypothetical unified world market, where the world price index $P_W = (\sum_i (\beta_i p_i \Pi_i)^{1-\sigma})^{1/(1-\sigma)}$ is solved from summing Equation 8. $P_W = 1$ is a convenient normalization of this hypothetical world price. Then with given $\beta_i p_i$'s, the normalization $(\sum_i (\beta_i p_i \Pi_i)^{1-\sigma})^{1/(1-\sigma)} = 1$ is a useful normalization in solving for multilateral resistances with Equations 6 and 7. The factor Π_i is straightforwardly interpreted as the sellers' incidence of trade costs from origin i (Anderson & Yotov 2010a).

¹For any solution to the system $\{P_j^0, \Pi_i^0\}$, $\{\lambda P_j^0, \Pi_i^0/\lambda\}$ is also a solution. Thus a normalization is needed. Anderson & Yotov (2010a) find that the system given in Equations 6 and 7 solves quite quickly, not surprisingly because it is quadratic in the $1 - \sigma$ power transforms of the P 's and Π 's.

P_j is now interpreted as buyers' incidence. Solving Equation 7 for P_j , it is a CES index of bilateral buyers' incidences $t_{ij}/\Pi_i, \forall i$, equivalent to buyers paying a uniform markup factor P_j on their entire bundle of shipments (from all i). Sellers' incidence Π_i similarly is then interpretable as a CES index of the bilateral sellers' incidences t_{ij}/P_j , from Equation 6.

The interpretation of Π and P as buyers' and sellers' incidence generalizes the elementary economics idea of incidence in the one-good case. If the actual set of trade costs were to be replaced with hypothetical trade costs $\tilde{t}_{ij} = \Pi_i P_j$, market clearance, and budget constraints, Equations 6 and 7 would still hold with the initial equilibrium shares; hence the sellers' factory gate prices would remain the same, and the aggregate buyers' prices would remain the same.² In this sense, the set of bilateral t_{ij} is equivalent to the set of \tilde{t}_{ij} 's that decompose into the product of buyers' and sellers' incidence factors. (Unlike the one-good case, it is the aggregate sales and purchases that is constant; bilateral flows would change in the hypothetical equilibrium.)

The model given in Equations 5–7 applies to any goods class. For disaggregated gravity, all variables and parameters given in Equations 5–7 should be understood as having superscript k 's to denote the goods class in question. When accounting for substitution between goods classes, aggregate expenditure (or the cost of intermediate inputs) is given by the expenditure (or cost) function $C(P_j^1, \dots, P_j^K)u_j$, where $C(\cdot)$ is the aggregate cost of living index for j , and u_j is the utility of the representative agent (or quantity of aggregate output). Then, by Shephard's lemma, $E_j^k = P_j^k \partial C(\cdot) / \partial P_j^k$. Each class of goods has expenditure shares described by Equations 3 and 4 but is amended with the addition of a superscript k to every variable and parameter. Anderson & van Wincoop (2004) argue theoretically for estimating disaggregated gravity, and Anderson & Yotov (2010a, 2010b) demonstrate that aggregation bias is large in practice.

The buyers' and sellers' incidence measures are usefully interpreted as the incidence of total factor productivity (TFP) frictions in distribution. They contrast with standard TFP-type measures of productivity in distribution. The sectoral TFP friction in distribution is defined by the uniform friction that preserves the value of sectoral shipments at destination prices: $\bar{t}_i^k = \sum_j t_{ij}^k y_{ij}^k / \sum_j y_{ij}^k$, where y_{ij}^k denotes the number of units of product class k received from i at destination j . \bar{t}_i^k is a Laspeyres index of outward trade frictions facing seller i in good k .

The TFP measure \bar{t}_i^k is useful for analyzing the distribution productivity of the world economy as a whole, but it is misleading for purposes of understanding comparative economic performance and the national patterns of production and trade. \bar{t}_i^k gives the sellers' incidence only under the partial equilibrium and inconsistent assumption that all incidence falls on the seller i . Anderson & Yotov (2010a,b) show that in practice these differences are significant: Laspeyres TFP measures and the incidence of TFP in distribution differ in magnitude, and in the case of inward measures, the correlation between them is low.³

²This property of Equations 6 and 7 was noted by Anderson & van Wincoop (2004), foreshadowing the interpretation of multilateral resistance as incidence.

³An alternative measure proposed by Redding & Venables (2004) resembles multilateral resistance but does not measure incidence. Their measure of market access uses essentially the same formula as Equation 6, whereas their measure of supplier access uses the CES price index formula $P_j^k = \left[\sum_i (\beta_i^k P_i^k t_{ij}^k)^{1-\sigma_k} \right]^{1/(1-\sigma_k)}$. These variables are constructed without taking account of the simultaneous determination of the two variables, so they do not measure incidence.

For consistency of the gravity modules with full general equilibrium, involving allocation across the sectors k in each country, the Π 's are normalized in each sector k for given parameters and factory gate price p_i^k by

$$\sum_j (\beta_i^k p_i^k \Pi_i^k)^{1-\sigma_k} = 1. \quad (9)$$

In practice, when analyzing a gravity module, it is often convenient to normalize one of the P 's to 1. The choice of normalization is irrelevant to the distribution of the goods because only relative incidence matters.

Now let us return to the interpretation of the gravity equation (Equation 5), reproduced here for convenience:

$$X_{ij} = \frac{Y_i E_j}{Y} \left(\frac{t_{ij}}{\Pi_i P_j} \right)^{1-\sigma}.$$

The right-hand side is the product of two ratios. The first ratio is the predicted frictionless trade flow given the E 's and Y 's, $Y_i E_j / Y$. The second ratio is thus interpreted as the ratio of predicted trade (given the t 's) to predicted frictionless trade.

The useful measure of constructed home bias (CHB) (Anderson & Yotov 2010a) is interpreted as the predicted value of internal trade of i with itself to the predicted value of internal trade in the frictionless equilibrium. CHB is thus given by

$$CHB_i \equiv \left(\frac{t_{ii}}{\Pi_i P_i} \right)^{1-\sigma}. \quad (10)$$

CHB varies substantially by country, product, and time because of changing expenditure and supply shares, even when gravity coefficients are constant (Anderson & Yotov 2010a,b).

Policy makers often focus on overall import penetration ratios such as $\sum_{i \neq j} X_{ij} / E_j$ and the analogous ratio $\sum_{i \neq j} X_{ji} / Y_j$ for exports. These concerns are acute for certain goods classes. The import and export penetration ratios are a linear function of CHB for any goods class k :

$$\sum_{i \neq j} X_{ij}^k / E_j^k = 1 - (t_{jj}^k / P_j^k \Pi_j^k)^{1-\sigma_k} Y_j^k / Y^k, \quad (11)$$

$$\sum_{i \neq j} X_{ji}^k / Y_j^k = 1 - (t_{jj}^k / P_j^k \Pi_j^k)^{1-\sigma_k} E_j^k / Y^k. \quad (12)$$

Anderson & Yotov show that CHBs vary a lot across goods, and more importantly for policy concerns, they exhibit a lot of intertemporal movement because of changing world shipment shares at constant t_{ij}^k 's, implying a lot of explanatory power over the import and export ratios.

The interpretation of the second ratio in Equation 5 applies straightforwardly to any bilateral flow: It is equal to the ratio of predicted bilateral trade to predicted frictionless trade; hence $(t_{ij} / \Pi_i P_j)^{1-\sigma}$ is the constructed trade bias on the link from i to j because of the buyer's bilateral incidence from i relative to the average buyer's incidence for country j . Alternatively, the same statistic viewed from the exporter's viewpoint results from the bilateral seller's incidence relative to the average seller's incidence. Bilateral trade flows

shift about because of changes in production and expenditure shares of world shipments, as implied by the frictionless gravity model, but also because of the general equilibrium force of share changes that alters incidence even when trade costs $\{t_{ij}\}$ are constant (Anderson & Yotov 2010a).

The gravity model also readily disaggregates within countries, allowing useful investigations of interregional versus international trade costs. Indeed, the development of the structural gravity model (Anderson & van Wincoop 2003) was provoked to solve a puzzle posed by one of the most provocative and useful empirical findings of the traditional gravity literature. McCallum (1995) found that crossing the Canadian border had an enormous trade-destroying effect on the trade flows of Canada's provinces. He found that Canada's provinces traded 22 times more with each other than with U.S. states, all else equal. This was too large to be a sensible component of bilateral trade costs t_{ij} .

Structural gravity solved the puzzle by showing that the border dummy variable in McCallum's traditional model reflected the effect of multilateral resistance. The border dummy in the McCallum regression shifts the ratio of interprovincial trade to province-state trade. Because it is a traditional gravity regression, it does not control for multilateral resistance. Using Equation 5 to form this ratio for a pair of such flows in the structural gravity model and rearranging terms, one obtains, for British Columbia's exports to adjacent Alberta and across the U.S. border to adjacent Washington,

$$\frac{X_{BC,AB}}{X_{BC,WA}} = \left(\frac{t_{BC,WA} P_{AB}}{t_{BC,AB} P_{WA}} \right)^{\sigma-1}.$$

The expression on the right-hand side of the equation reflects not only the direct trade cost increase at the U.S. border that raises $t_{BC,WA}/t_{BC,AB}$, but also the effect of the ratio of multilateral resistances for a province and a state, in this case Alberta and Washington, P_{AB}/P_{WA} . Because Canada's provinces must do far more of their trade with the outside world than do U.S. states (Canada is about one-tenth the size of the United States in GDP), the provinces naturally have higher multilateral resistance than the states, thereby greatly increasing interprovincial trade. In McCallum's traditional gravity regression, the border dummy variable has a regression coefficient that is an average of such terms, although a biased estimate of it because of the omission of the multilateral resistance controls from his regression. Estimating the structural gravity model, Anderson & van Wincoop (2003) find a more plausible border-cost component of t_{ij} , in the range of 20%–50%.

Interregional versus international trade cost implications of structural gravity were further developed by Anderson & Yotov (2010a). They offer a decomposition of incidence into domestic and international components and calculate sellers' incidence for Canada's provinces on trade within Canada as compared with trade with the rest of the world. They find that, although incidence overall declined substantially from 1990 to 2002, it was entirely on the external trade; sellers' incidence on domestic trade remained constant. Similar investigations are likely to provide a useful context for regional integration policy in many countries and economic areas around the world where separatism and economic integration are important concerns.

Notice that the trade flows in Equation 5 are invariant to a uniform rise in trade costs (including costs of internal shipment). This follows because Equations 6 and 7 imply that raising all t_{ij} 's by the factor $\lambda > 1$ will raise each Π and P by the factor $\lambda^{1/2}$. This formal homogeneity property has useful empirical content: If the world really were getting smaller

uniformly, the gravity model would be unable to reveal it. The empirical literature tends to indicate little change in gravity coefficients (see, especially, Anderson & Yotov 2010a,b), contrary to intuition about globalization driven by falling communications costs and improving quality of transport but consistent with uniform shrinkage of resistance to trade.

Anderson (1979) was the first to derive gravity from the Armington/CES preference structure, noting that Armington preferences implied a bilateral trade flow gravity equation of the form of Equation 5 that would require controlling for the importer and exporter trade cost indexes. By using a units choice to set all equilibrium factory gate prices equal to 1, Anderson's (1979) derivation concealed how Equations 5–7 formed a conditional general equilibrium module that would be the foundation for the useful comparative statics to come a generation later. The comparative statics of inward and outward multilateral resistance were first used by Anderson & van Wincoop (2003). Recognition that multilateral resistance is interpreted as incidence is in Anderson & Yotov (2010a).

4.2. Supply-Side Structure

An alternative derivation of a mathematically equivalent structural gravity model was proposed by Eaton & Kortum (2002), based on homogeneous goods on the demand side, iceberg trade costs, and Ricardian technology with heterogeneous productivity for each country and good due to random productivity draws from a Fréchet distribution. Despite CES structure for the intermediate goods demand, in equilibrium the share of goods demanded from i by country j is determined only on the supply side; the influence of σ disappears into a constant term. In equilibrium each country will be assigned a subset of the goods, and except for knife-edge cases, it is the only supplier of these goods. The bilateral trade flows obey the same formulas as Equations 5–7. $1 - \sigma$ is interpreted as $-\theta$, where θ is the dispersion parameter of the Fréchet distribution. In contrast to the Armington/CES model, all action is on the extensive margin of trade. Eaton & Kortum derive their model for one sector only, a specification generalized by Costinot et al. (2010), so that θ_k is the dispersion parameter for the distribution describing productivity draws in sector k .

The Ricardian structure of supply leads to a simple general equilibrium superstructure, an appealing feature that has led to a growing literature combining estimation and simulation. The general equilibrium superstructure is discussed below in Section 7.

Chaney (2008) derives a similar supply-side gravity structure based on Ricardian productivity draws from a Pareto distribution in which the dispersion parameter of the Pareto distribution plays essentially the same role as θ in the Eaton-Kortum model. For each firm, changes in variable trade costs act on the intensive margin, but for the total sectoral bilateral trade flow, these effects disappear and the aggregate effect is effectively on the extensive margin of trade. Chaney's model includes a fixed cost of export for monopolistically competitive firms, and in equilibrium the elasticity of substitution affects the pattern of trade by being part of the elasticity of equilibrium trade volume with respect to the fixed cost.

4.3. Zeroes

In practice, many potential bilateral trade flows are not active. The data presented to the analyst may record a zero that is a true zero or it may reflect shipments that fall below a

threshold above zero. In addition there may be missing observations that may or may not reflect true zeroes. The prevalence of zeroes rises with disaggregation, so that in finely grained data, a large majority of bilateral flows appear to be inactive. Finally, over time, the small bilateral flows in finely disaggregated data appear to wink on and off. The zeroes present two distinct issues for the analyst: appropriate specification of the economic model and appropriate specification of the error term on which to base econometric inference. Discussion of the specification of the error term is deferred to Section 5.

In specifying the economic model, zero trade flows present a problem for the CES/Armington model of demand and the Eaton-Kortum supply-side structure. With elasticities of substitution greater than one (or the equivalent dispersion/comparative advantage parameter restriction for the Eaton-Kortum model), the empirically relevant case, some volume will be purchased no matter how high the price. One way to rationalize zeroes is to modify the demand specification so as to allow choke prices above which all demand is choked off. A start is made by Novy (2010), who derives gravity in a highly restricted one-slope parameter translog expenditure function case that allows for zeroes in demand.⁴ More general translog treatments are feasible and desirable. Anderson & Neary (2005) present a general homothetic preferences structure, showing that multilateral resistance is defined and solved from a similar equation system once the functional form and its parameters are specified, along with data on shipment and expenditure shares.

An alternative economic specification explanation retains CES/Armington preferences and rationalizes zeroes as resulting from fixed costs of export facing monopolistic competitive firms. If no firm in i is productive enough to make incurring the fixed cost of exporting to j profitable (given the cost of production in i , variable trade cost t_{ij} , and willingness to pay in j), then zero trade results. Helpman, Melitz, and Rubinstein (HMR; Helpman et al. 2008) develop this idea. The selection effect determines which markets are active and also determines a volume effect V_{ij} due to productivity heterogeneity among firms whereby markets that are active have a greater or lesser number of firms active depending on the same selection mechanism. The gravity model becomes

$$X_{ij}^k = \frac{E_j^k Y_i^k}{Y^k} V_{ij}^k \left(\frac{t_{ij}^k}{P_j^k \Pi_i^k} \right)^{1-\sigma_k},$$

$$(\Pi_i^k)^{1-\sigma_k} = \sum_j \left(\frac{t_{ij}^k}{P_j^k} \right)^{1-\sigma_k} \frac{V_{ij}^k E_j^k}{Y^k},$$

$$(P_j^k)^{1-\sigma_k} = \sum_i \left(\frac{t_{ij}^k}{\Pi_i^k} \right)^{1-\sigma_k} \frac{V_{ij}^k Y_i^k}{Y^k}.$$

HMR report results suggesting that this mechanism is indeed potent and that inference without accounting for it biases estimates of the variable trade costs downward.

The key mechanism is a Pareto productivity distribution of potential trading firms. The Pareto distribution is capable of capturing the empirical observation that the largest and most productive firms export the most and to the most destinations. The Pareto distribution allows a tractable estimation procedure that requires only aggregate bilateral trade

⁴Novy's aggregate bilateral OECD trade flow data contain no zeroes, so this feature is not exploited yet.

data, an important advantage because firm-level trade data are not widely available. In practice, identification of the parameters in estimating the HMR model requires a plausible exclusion restriction—a proxy for the fixed cost of export that is not also a proxy for the variable cost of trade. HMR use common religion, a specification that many find dubious.

An important challenge for the future is combining the HMR mechanism with the translog expenditure system. This combination might be able to distinguish between fixed export costs and choke prices as an explanation of zeroes.

4.4. Discrete Choice Structure

The third alternative model of structural gravity is based on modeling individual discrete choice in a setting in which the individual trader faces costs or receives benefits not observable to the econometrician. Of all possible bilateral pairs, the trader chooses one because it yields the greatest gain. A population of such traders has observable characteristics such as bilateral distance that condition the probability of each choice; the econometrician observes the resulting masses allocated and uses a probability model to structure statistical inference. An early attempt along these lines was made by Savage & Deutsch (1960) and followed by Leamer & Stern (1970). Several problems with the model limited its appeal. It did not offer a rationale for the linear homogeneity of the mass variables in gravity, and its characterization of cross effects did not have a sound rationale.

Discrete choice modeling was greatly advanced by McFadden (1974), who proved that under plausible restrictions in this setting (the random variable, to the econometrician, results in the observed choices following the type 1 extreme value distribution), the resulting probability model is the multinomial logit. Building on the multinomial logit, it is easy to generate a structural gravity model. This reasoning has rationalized recent work on models of migration (e.g., Grogger & Hanson 2008, Beine et al. 2009).

It is straightforward to combine the discrete choice setup with the market clearance conditions to derive the buyers' and sellers' incidence of trade costs exactly as in the preceding models. The development is postponed to the next section, but is noted here because exactly the same reasoning applies to goods traders making discrete choices on where to sell or buy their goods. Thus the discrete choice probability model rationalizes structural gravity equally well. It may be fruitful to explore the applicability of two-sided matching models in the trade context as well as the job-market context.

5. ESTIMATION

As an empirical model, gravity is fundamentally about inferring trade costs in a setting in which much of what impedes trade is not observable to the econometrician. What are observable are the trade flows and a set of proxies for various types of trade costs, along with direct measures of some components of trade costs. Most issues with modeling trade costs are discussed in Anderson & van Wincoop (2004). Since that time, there have been several notable advances in modeling and inferring trade costs.

Two of the advances deal with the implications of zeroes in the bilateral trade flow data. One view of zeroes is that they stand for flows too small to report, an interpretation that indeed represents reporting practices of government trade ministries. Interpreting zeroes in this way, it is legitimate to drop the zero observations from estimation because there is no economic significance to the zeroes relative to the nonzero observations.

In the presence of heteroskedastic errors, Santos Silva & Tenreyro (2006) point out that inconsistent estimation arises from the usual econometric gravity practice using logarithmic transforms of Equation 5 augmented with a normal disturbance term and estimated with ordinary least squares (OLS). Because the data have a lot of zeroes, the disturbance term must have a substantial mass at very small values, violating the normal distribution assumption. They propose instead to model the disturbance term as generated from a Poisson distribution, leading to estimation with a Poisson pseudo-maximum likelihood (PPML) technique. Their results show that PPML leads to smaller estimates of trade costs compared to OLS.

The heteroskedastic error problem identified by Santos Silva and Tenreyro is important, but their solution has not convinced all researchers. Martin & Pham (2008) argue, based on Monte Carlo simulations, that when heteroskedasticity is properly controlled, Tobit estimators outperform PPML when zeroes are common. Heteroskedasticity is likely to be attenuated using size-adjusted trade X_{ij}/Y_iE_j as the dependent variable, as advocated by Anderson & van Wincoop (2003, 2004).

An alternative view of zeroes, encountered above, is that economically meaningful selection generates the zeroes. All firms in origin i face fixed costs of entering exporting to any particular destination j , and only the sufficiently productive ones can afford to pay the fixed cost. When a destination j is so expensive to reach that no firm in i can afford the fixed cost, zeroes are generated in the data. In this case, OLS estimation without accounting for selection is biased for two reasons: the standard left censored selection reason and because, for bilateral pairs with positive flows, of a volume effect due to selection of firms along with the bilateral trade cost t_{ij} that is the object of investigation in OLS or PPML estimation. HMR find that their technique also results in lower cost estimates than does OLS. (They report that estimation with Poisson error terms as opposed to normal ones does not alter their findings.)

In principle, an economic model of zeroes is attractive, but many researchers are suspicious of the exclusion restriction used by HMR to identify their volume effect. They assume that common religion affects the fixed cost of export but not the variable cost t_{ij} . Moreover, the tractability of the HMR model depends on a restrictive distributional assumption on the productivity draws distribution of firms, which in turn is a specialization of a particular model of monopolistic competition that is not applicable to all sectors.

Anderson & Yotov (2010b) report that estimation with PPML, HMR, or OLS leads to essentially identical results for buyers' and sellers' resistance and CHB because it leads to gravity coefficients that are almost perfectly correlated. The homogeneity property of Equations 5–7 implies that only relative trade costs can be inferred by gravity; hence the differences in techniques effectively amount to different implicit normalizations. Anderson & Yotov report this near-perfect correlation finding based on estimation with the three techniques over 18 three-digit manufacturing sectors, 76 countries, and 13 years of data.

5.1. Traditional

Some researchers continue to use a traditional form of the gravity model, presumably in the belief that the structural model featured above is not sufficiently well established. It seems useful to review a generic traditional model along with my objections.

A typical traditional gravity model regresses the log of bilateral trade on log trade costs proxied by a vector of bilateral variables that are not at issue here, log GDP for origin and

destination, and log population for origin and destination. In addition, a number of authors include remoteness indexes of each country's distance from its partners, atheoretic measures that are inadequate attempts to control for multilateral resistance. [Anderson & van Wincoop (2003) report significant differences between gravity estimated with remoteness and with multilateral resistance.]

The first objection to the traditional model is its aggregation, which causes two problems. There is aggregation bias because of sectorally varying trade costs and sectorally varying elasticities of trade with respect to costs (see Anderson & van Wincoop 2004 for analysis and Anderson & Yotov 2010a,b for evidence on downward bias). The second aggregation problem is specification bias because GDP is a value-added concept with a variable relationship to gross trade flows. Much recent attention to the vertical disintegration of production and its international aspect emphasizes the variable intertemporal relationship of gross trade to GDP, and its variation across countries is also significant. Disaggregation and use of the appropriate sectoral output and expenditure variables fix both problems.

The second objection is omitted variable bias from the perspective of the structural gravity model—the traditional model leaves out multilateral resistance. Multilateral resistance has only low correlation with remoteness indexes, and the omitted variable will be correlated with the other right-hand-side variables and thus bias estimation. The traditional model's inclusion of mass variables such as GDP and population presumably picks up a part of the missing explanatory power of multilateral resistance, as Anderson and Yotov's work shows that multilateral resistance is associated with country size. Estimation with country fixed effects controls appropriately for all these issues.

5.2. Structural

Anderson & van Wincoop (2003) combine Equations 6 and 7 with the stochastic version of Equation 5 to form a full information estimator of the coefficients of the proxies for trade costs such as distance and international borders. Utilizing the unitary elasticities on the E 's and Y 's, their dependent variable is $X_{ij}/Y_i E_j$, size-adjusted trade.

An alternative fixed effects estimator controls for the unobservable multilateral resistances and activity variables:

$$X_{ij} = x_i m_j t_{ij}^{1-\sigma} \varepsilon_{ij}, \quad (13)$$

where ε_{ij} is the random error term, x_i is the fixed effect for country i as an exporter, and m_j is the fixed effect for country j as an importer. Equation 13 is less efficient than a full information estimator but seems preferable to most subsequent investigators. Feenstra (2004) argues for the fixed effects estimator because it does not require custom coding; however, another and perhaps better reason is that researchers should be suspicious that there may be other country-specific unobservables that the fixed effects pick up, but that full information estimation would drive toward spurious results.

A major drawback to fixed effects estimation is its demolition of structure: The econometrician blows up the building to get at the safe inside containing the inferred bilateral trade costs. Fortunately, in the case of structural gravity, it is feasible to reconstruct the building like an archeologist, using structural principles in the form of Equations 6 and 7. Thus Anderson & Yotov (2010a,b) use fixed effects to estimate Equation 5 in its stochastic

form, but then estimate the multilateral resistances by calculating the fitted t_{ij} 's and plugging them into Equations 6 and 7.

This technique is used to test the structural gravity model by comparing the estimates of fixed effects ($\alpha_i m_j$) with the structural gravity term ($Y_i E_j \Pi_i^{\sigma-1} P_j^{\sigma-1}$). The results are remarkably close in an economic sense (the fitted regression line has an estimated elasticity around 0.96, compared with the theoretical value of 1.0) across 76 countries and 18 manufacturing sectors over 13 years. Although this result suggests that the constraints that legitimize full information methods are very close to being valid, fixed effects estimation still seems the better, more cautious practice to follow.

Baier & Bergstrand (2009) propose an alternative direct estimator of multilateral resistance based on a Taylor's series approximation of Equation 5. They report reasonably good results, but I suspect that many researchers will be wary of the approximation error. In contrast, the method of Anderson and Yotov avoids approximation error. As Baier & Bergstrand emphasize, the advantage of their method relative to panel estimation with fixed effects is that it avoids the upper bound on the number of fixed effects imposed by typical econometric packages at the time of this writing. (STATA currently imposes a limit of 11,000 independent variables, whereas 100 countries over 10 years require approximately 200,000 fixed effects, and even yearly estimation requires 20,000.) In principle, the Baier & Bergstrand estimator could be used to construct t_{ij} 's and then could be combined with data on the Y 's and E 's using Equations 6 and 7 to obtain the incidence measures and perform comparative statics with them. The constructed multilateral resistances in Baier & Bergstrand's method can be compared to the point estimates, differences being attributed to random error and approximation error.

5.3. Foreign Affiliate Sales

A large share of international trade comprises sales by foreign affiliates of multinational firms. Standard trade gravity models include this trade along with that of domestically owned firms. If the trade costs are the same for both types of firms, this treatment is entirely appropriate.

There is reason to believe, however, that the trade cost structure facing foreign affiliate sales differs from that facing domestic firms. For trade in intermediate inputs, information and other transactions costs are reduced for intrafirm trade, but even for horizontal trade there are likely to be transactions cost advantages when a foreign affiliate sells into its home country. This reasoning suggests that one should split the home and foreign firms into separate sectors for more accurate and informative inference about trade costs.

This approach to gravity with multinationals follows the conditional general equilibrium strategy, treating total sales as exogenous. It avoids taking a stand on determinants of the location of production. A significant literature that is at least loosely related to gravity attempts to explain this location decision along with the volume of foreign affiliate sales. It is treated below in the discussion of foreign direct investment (FDI).

6. GRAVITY AND FACTOR FLOWS

Gravity has long been applied to empirically model factor movements. As with trade flows, the model always fits well. But, in contrast to the recent development of an economic structural gravity model of trade, there has been little progress in building a theoretical

foundation. This section sets out a structural model of migration, reviews promising steps toward a structural model of FDI, and closes by pointing to the unsolved puzzle of modeling international portfolio capital movements.

6.1. Migration

The decision to migrate is a discrete choice from a menu of locations. Each worker that migrates faces a flow cost common to all workers who migrate in a particular bilateral link, but each worker also has an idiosyncratic component of cost or utility from the move. We may think of an idiosyncratic cost component as plausibly associated with a fixed cost, but in the migration decision, the distinction between fixed and variable cost plays no important role because the decision to migrate has no volume decision accompanying it. This stands in contrast to the export selection model of Helpman et al. (2008) in which the decision to export and the decision of how much to export are distinct.

Let w^i denote the wage at location $i, \forall i$. The worker b who migrates from origin j to destination i faces a cost of migration modeled with iceberg cost factor $\delta^{ij} > 1$, receiving net wage (w^i / δ^{ij}) . Worker b 's idiosyncratic utility from migration is represented by ϵ^{ijb} , private information to him. He chooses to migrate if $(w^i / \delta^{ij}) \epsilon^{ijb} \geq w^j$ for at least some i . Among alternative destinations, he chooses the one with the largest surplus. Suppose that the worker has logarithmic utility. Then his observable component of utility of migration from j to i is $u^{ij} = \ln w^i - \ln \delta^{ij} - \ln w^j$. In this sort of setting, McFadden (1974) shows that if ϵ had the type 1 extreme value distribution, the probability that a randomly drawn individual would pick any particular migration destination is given by the multinomial logit form.

Building on this insight, migration models subsequently used the multinomial logit to model bilateral migration flows (for two recent examples, see Beine et al. 2009, Grogger & Hanson 2008). This section develops a novel gravity model representation of the migration model by making use of the market-clearing conditions to derive the appropriate multilateral resistance variables.

At the aggregate level, the probability is equal to the proportion of migrants from j (assumed to be identical except for their values of ϵ) that pick destination i . Let N^j denote the population of natives of j . The predicted migration flow from j to i that results from the setup is

$$M^{ij} = G(u^{ij})N^j, \tag{14}$$

where

$$G(u^{ij}) = \frac{\exp(u^{ij})}{\sum_k \exp(u^{jk})}.$$

With logarithmic utility, the migration equation is

$$M^{ij} = \frac{w^i / \delta^{ij}}{\sum_k w^k / \delta^{jk}} N^j. \tag{15}$$

Equation 15 is a structure analogous to the CES demand (in the Armington model) or Ricardian supply (in the Eaton-Kortum model) shares that underpin the trade gravity equation. The connection of the share equation (Equation 15) to the structural gravity form of the model is completed by using the labor market balance equations to solve for and substitute out the equilibrium w 's.

In natural logarithm.

Define $W^i \equiv \sum_k w^k / \delta^{ik}$ and define the labor force supplied to i from all origins:

$$L^i \equiv \sum_j M^{ji}. \quad (16)$$

Also, $N \equiv \sum_j N^j = \sum_i L^i$, the world labor supply N . The labor market clearance equation is

$$L^i = w^i \sum_j ((1/\delta^{ji})/W^j)N^j.$$

Then

$$w^i = \frac{L^i}{\Omega^i N}, \quad (17)$$

where

$$\Omega^i = \sum_j \frac{1/\delta^{ji} N^j}{W^j N}. \quad (18)$$

Using Equation 17 to substitute for the wage in W^j ,

$$W^j = \sum_k \frac{1/\delta^{jk} L^k}{\Omega^k N}. \quad (19)$$

Substituting for the wage in Equation 15 using Equation 17 yields the structural gravity equation of migration:

$$M^{ji} = \frac{L^i N^j}{N} \frac{1/\delta^{ji}}{\Omega^i W^j}. \quad (20)$$

The first ratio represents the migration pattern of a frictionless world. The implication is that in a frictionless world, populations originating in j would be found in equal proportions to their share of world population in all destinations: $M^{ji}/L^i = N^j/N$. The second term represents the effect of migration frictions. The bilateral migration friction δ^{ji} reduces migration. It is divided by the product of the weighted averages of the inverse of migration frictions, one for inward migration to i from all origins and one for outward migration from j to all destinations. The system given in Equations 18 and 19 can be solved for the Ω 's and W 's (subject to a normalization). Their interpretation and their connection to multilateral resistance in the more familiar trade gravity model are easier to see in the case in which utility is generalized to the log of a constant relative risk aversion function.⁵

Let the coefficient of relative risk aversion be θ . In this case Equation 20 becomes

$$M^{ji} = \frac{L^i N^j}{N} \left(\frac{\delta^{ji}}{\Omega^i W^j} \right)^{1-\theta},$$

⁵A tractable gravity equation results from Equation 14 by using a restriction on utility to convert $\exp u^{ji}$ into a tractable form. When utility is given by the log of any power function of the wage net of migration costs, the CES-type form of gravity results, with consequent ease of estimation and resemblance to the trade flow structural gravity model.

where

$$\bar{\Omega}^i \equiv \left[\sum_j \frac{(\delta^{ji})^{1-\theta} N^j}{\bar{W}^j N} \right]^{1/(1-\theta)}$$

and

$$\bar{W}^j \equiv \left[\sum_i \frac{(\delta^{ij})^{1-\theta} L^i}{\bar{\Omega}^i N} \right]^{1/(1-\theta)}.$$

Here, $\bar{\Omega}^i$ and \bar{W}^j are CES price indexes of migration frictions, one for inward ($\bar{\Omega}^i$) and one for outward (\bar{W}^j) migration frictions. These equations are exactly analogous to the Anderson and van Wincoop model's inward and outward multilateral resistance equations for trade, but applied to migration. As with the trade gravity model, outward multilateral resistance gives the sellers' incidence of the migration costs on average while the inward multilateral resistance gives the buyers' incidence of migration costs. Equations 18–20 result from the special case $\theta = 2$.

Ω and W are general equilibrium concepts as is clear because their solution in the simultaneous systems above involves every bilateral migration cost in the world. They are conditional general equilibrium concepts because the L 's are endogenous in a full general equilibrium. It is possible in a Ricardian production setting to combine the migration system with the trade gravity model to derive equilibrium labor supplies that are functions of the incidence of both migration frictions and trade frictions.

As with trade gravity models, $\bar{\Omega}$'s and \bar{W} 's can be computed once the δ 's are econometrically constructed and the labor supplies L^i and population stocks N^j are observed. A normalization is needed (see Anderson & Yotov 2010a for details).

A similar model has been applied to services trade by Head et al. (2009). Instead of actually changing locations, the foreign worker does the job in her home location. The cost of migration becomes the cost of monitoring the distant worker. Worker productivities in each location have the Fréchet distribution, as in the Eaton-Kortum model. The firm selects workers so as to minimize the log of the delivered unit labor cost. Then the distribution of log productivities takes the Gumbel form. The fraction of service jobs in origin i going to workers in location j has the multinomial logit form. The total number of workers and of jobs in each location enters the model in the same way as in the migration model. I suspect that the choice between offshoring the service job and migrating the worker can be fruitfully addressed with some combination of the two models.

The preceding treatment applies to a stationary equilibrium in which the L 's are the result of M 's fully adjusting labor supplied at each location to its equilibrium value given the initial stocks of labor $\{N^j\}$ and the set of migration frictions, the δ 's. In adapting the model to fit actual data, the N 's, L 's, and M 's are observed at points in time, and with panel data, the observations are linked over time.

If the sequence of observations is regarded as reaching the static equilibrium each period, the observed migration is just that amount needed to reach the equilibrium in each period. This model would be consistent with naive expectations about future wages, or with a pure guest worker model in which migration is determined by contemporaneous variables only. So, in principle, under this interpretation, the preceding model could be applied at each date, all variables now having a time subscript.

The alternative is a dynamic model in which the migrants form expectations about the sequence of future wages based on underlying expectations about the future evolution of the distribution of trade frictions, the populations, and, as shown below following the development of the integrated trade and migration model, variables that predict the demand for labor at all locations. This sophistication requires a big increase in complexity, with dubious applicability of rational expectations to unskilled workers.

The other issue raised by thinking of dynamics is the issue of partial adjustment—migration in any one year may not suffice to reach the static equilibrium of the preceding section. In this case, the standard ad hoc approach of partial adjustment due to quadratic adjustment costs might be applied without too large an increase in complexity.

6.2. Foreign Direct Investment

FDI has been successfully explained by gravity structures without a theoretical foundation. More recent work has made progress on foundations. Satisfactory foundations are more difficult to find for at least two reasons. First, the question of the location of production must be answered in an upper-level general equilibrium model, which requires taking a stand on one of many possible production, preference, and market structures restricted so as to produce tractable results. Second, the determinants of location depend on whether the good in question is vertically or horizontally linked to other sources of firm profits.

A key element in explaining the location of horizontally linked production is the proximity-concentration tradeoff: A firm with fixed cost reduces per-unit production cost by concentrating production at one location but can save distribution costs by allocating production in proximity to markets. Even under strong restrictions, the models obtained so far are nonlinear and require approximation to be taken to data.

Helpman et al. (2004) model interaction between horizontally linked exports and foreign affiliate sales, in which the firm chooses between exporting from home or investing abroad and selling from a foreign plant. They are able to draw inferences from aggregate data by modeling heterogeneous productivity of firms with a Pareto distribution. Fixed costs of export and of investing abroad serve to select firms into nontraders, exporters, and multinationals with ratios that vary market by market because of trade costs modeled as transport costs and tariffs only, omitting the usual gravity variables. Their empirical application with U.S. data obtains fairly good results in explaining the ratio of exports to foreign affiliate sales with a linear approximation to their underlying nonlinear model. The model fits much less well than standard export gravity equations, which is not surprising because the dependent variable is different and the question addressed is more difficult to answer.

Kleinert & Toubal (2010) extend Helpman et al. (2004) to allow for fixed setup costs that rise with distance, a wrinkle that can explain why foreign affiliate sales can fall rather than rise with distance, as the earlier proximity-concentration tradeoff suggested. They also derive a gravity-type relationship from two other structures, a vertical integration model and a two-country factor proportions model of fragmentation.

Bergstrand & Egger (2007) offer a gravity model of FDI derived from the knowledge-based capital theory of horizontal multinational enterprises. Their objective is a full general equilibrium model that can explain trade, foreign affiliate sales, and FDI. They simulate a theoretical model that generates nonlinear relationships between exports, affiliate sales, and their exogenous determinants. Then they fit an approximate empirical

relationship to the generated data and take the same relationship to actual data, with some success. A limitation of their model is that, although the factor proportions model with three factors is used to explain simultaneous exports and affiliate sales, the countries in their simulation setup have identical endowment proportions and differ only in size.

Keller & Yeaple (2009) develop a gravity model of vertically integrated intrafirm trade featuring trade costs with two elements, a standard iceberg trade cost and a communication cost that rises with the complexity of the firm's technology. Input complexity raises technology transfer costs, whereas the costs of embodied technology transfer are independent of complexity and are increasing in trade costs. An increase in trade costs reduces foreign affiliate sales, and this effect is strongest in the most complex sectors. In contrast, an increase in trade costs reduces the imports of foreign affiliates, and this effect is weakest in the most complex sectors. Like the standard trade gravity model, Keller & Yeaple's model of foreign affiliate sales permits inference about trade costs from observable trade flows.

Keller & Yeaple report fairly good results estimating the model using confidential data on U.S. multinational firm activity from the Bureau of Economic Analysis. The role played by communication cost interacting with technological complexity thus appears likely to be helpful in explaining the rising share (in total trade) of intrafirm trade and also the rising share of trade in intermediates.

An alternative strategy along the lines of the conditional general equilibrium approach outlined above for migration appears useful. The migration decision model of Section 2 could apply to FDI because the location decision for a plant is similar to the location decision of a migrant. (Unlike migration but like trade, FDI involves a volume decision along with a participation decision.) The rate of return on investment could be taken as exogenous in a conditional general equilibrium approach just as wages are taken as exogenous in the migration gravity model, while market-clearing conditions apply just as in the migration model. Idiosyncratic cost factors would apply to the various investment projects, just as they do to the individual migrants. The Keller-Yeaple model of vertically integrated intrafirm trade offers a structure for identifying one type of cost. A weakness in the extension by analogy is in risk diversification. Migrants cannot diversify their risks, but firms can, although with limited possibilities that may be very limited for FDI. The potential risk diversification would modify the utility derived from each location choice. The discrete choice approach faces truly formidable modeling challenges in endogenizing the investment rates of return, unlike the wage equation suggested by the migration model.

A promising start on these lines is by Head & Ries (2008). Potential acquisitions go to the highest bidder, who bids based on his anticipated return net of monitoring costs that rise with distance and other standard gravity variables. The probability of the winning bid going to source country i takes the multinomial logit form. The mass variables are the stocks of projects in each host country and each source country's share of world bidders.

6.3. Portfolio Investment

Martin & Rey (2004) offer the first gravity-type model of international portfolio investment. The coefficient of relative risk aversion plays the role, in equilibrium, of the elasticity of substitution in the CES demand specification. Although appealing as a rationale for the gravity application of Portes & Rey (2005), the Martin & Rey model does not provide a fully satisfactory foundation for gravity models of investment flows because (a) trade is

assumed to be frictionless, (b) investment costs are uniform, and, most importantly, (c) the analysis is restricted to two countries. The third-party effects that play a big role in the gravity model of trade (and of migration) cannot be treated.

7. INTEGRATED SUPERSTRUCTURE

The gravity model nests inside a general equilibrium superstructure. As pointed out by Anderson & van Wincoop (2004), modularity implies that the problem of resource and expenditure allocation across sectors in the general equilibrium superstructure can be treated separably from the gravity module problem of distribution within sectors to destinations or from origins. Consistency between the two levels of the problem requires fixed-point calculations in general, but the economy of thought and computation due to separability is extremely useful and, in particular, makes it possible to integrate gravity with a wide class of general equilibrium production models. So far, only simple production models have been used for full general equilibrium comparative statics, but I anticipate that this situation will change.

The simplest production structure is an endowments economy. Anderson & van Wincoop (2003) use the endowments model to calculate the effect of eradicating the U.S.-Canada border on their estimated gravity model of trade between U.S. states, Canadian provinces, and the aggregated rest of the world.

Another attractive candidate is the Ricardian production model. Eaton & Kortum (2002) nest gravity inside a Ricardian model of production, a choice followed by a host of subsequent researchers such as Arkolakis (2008). An important feature of these models is the action on the extensive margin, as industries arise or disappear. In the 2002 Eaton-Kortum model, the extensive margin is the only margin. Arkolakis and others have variants in which both extensive and intensive margins are active. This is an important feature because disaggregated trade data and especially firm-level data indicate that both margins are active.

Between the two extremes of zero and infinite elasticity of transformation of the endowments and Ricardian models lie a host of more complex production structures in which action occurs on the intensive margin of production when relative prices change, leading to another channel of interaction between the gravity modules in each sector (and resulting buyers' and sellers' incidences) and the pattern of production. Consistency between the modules is achieved by using Equation 9 to normalize the Π 's in each sector. I think the future will see work with these more complex general equilibrium features.

Migration of labor and capital in the form of FDI has been given a complete gravity representation in this review. In the integrated superstructure, it can be treated simultaneously with the trade modules. In this setting, multilateral resistance in trade has significant effects on migration and vice versa. I anticipate that development of this link will be useful.

A number of authors have constructed integrated models that motivate econometric work aimed at discriminating between one or another specification of the upper-level production and market structure. A summary of work along these lines is presented in Feenstra (2004, chapter 5), where the main focus is on the link between gravity and increasing returns to scale. Research has continued on these lines, but I do not review it here.

I think the gravity model is a poor vehicle for inferences about returns to scale, market structure, and the global general equilibrium links between economies. This review and my previous work argue that gravity is about the distribution of given amounts of goods in each origin drawn by given amounts of expenditure in each destination, enabling inference about trade costs from the deviation of observed distribution from the frictionless equilibrium. The determinants of total shipments and total expenditures are irrelevant to this inference because country fixed effects are a consistent control that does not require taking a stand on any particular production or market structure model. Conversely, the cross-section variation of bilateral trade does not seem likely to have much useful information about the determination of national total shipments or expenditure. Interdependence is so deeply wound between these variables in the full general equilibrium model that inference about structure seems implausible. In contrast, simulation models look reasonably promising as a source of insight.

8. CONCLUSION

This idiosyncratic review of work on the gravity model suggests that the story is not over, so a conclusion can only point to potential future chapters. Distribution broadly defined consumes a very large share of the world's resources, and gravity has proven to be the most generally useful empirical model for understanding the distribution of goods and factors of production. It appears to work well at almost any scale.

The progress in structural modeling of gravity has yielded three distinct rationales for the same observationally equivalent model of the distribution of economic flows between origins and destinations, one based on the demand side (the CES/Armington model), one based on the supply side (the Eaton-Kortum model), and one based on a discrete choice model of the individual actor transferring the goods or factors. Further work may suggest ways to discriminate between these.

The structural modeling of gravity imposes trade separability, permitting gravity modules to be nested inside a wide range of general equilibrium superstructures. Future work with simulation models may suggest which of many candidate general equilibrium production models do better.

The problem of zeroes in the trade and factor flows data has been addressed with some success, particularly by HMR. But I expect future work to do better. The CES framework (with elasticity of substitution greater than one) is unsuitable for describing small amounts of trade. The translog cost function, in particular, seems likely to yield better descriptions and better understanding of why so many potential flows are equal to zero. This is so even if, as in HMR, fixed export costs play an important role in selecting firms to export.

Featured in this review are incidence measures produced by Anderson and Yotov. If the profession agrees that they are as interesting and useful as they appear to me, more work is needed to see how believable the measures are. As it stands, they are completely reliant on CES structure. How well does the CES do in representing the world economy? This is an especially important question in light of the zeroes question in the preceding paragraph. I look forward to the development of the translog case to help answer this question.

DISCLOSURE STATEMENT

The author is not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

I thank Jeffrey Bergstrand, Keith Head, and Yoto Yotov for helpful comments on earlier versions of this review.

LITERATURE CITED

- Anderson JE. 1979. A theoretical foundation for the gravity equation. *Am. Econ. Rev.* 69:106–16
- Anderson JE, Neary JP. 2005. *Measuring the Restrictiveness of International Trade Policy*. Cambridge: MIT Press
- Anderson JE, van Wincoop E. 2003. Gravity with gravitas. *Am. Econ. Rev.* 93:170–92
- Anderson JE, van Wincoop E. 2004. Trade costs. *J. Econ. Lit.* 42:691–751
- Anderson JE, Yotov YV. 2010a. The changing incidence of geography. *Am. Econ. Rev.* 100:2157–86
- Anderson JE, Yotov YV. 2010b. *Specialization: pro- and anti-globalizing, 1990–2002*. NBER Work. Pap. 16301
- Arkolakis C. 2008. *Market penetration costs and the new consumers margin in international trade*. NBER Work. Pap. 14214
- Baier SL, Bergstrand JH. 2001. The growth of world trade: tariffs, transport costs, and income similarity. *J. Int. Econ.* 53(1):1–27
- Baier SL, Bergstrand JH. 2009. Bonus vetus OLS: a simple method for approximating international trade-cost effects using the gravity equation. *J. Int. Econ.* 77(1):77–85
- Beine M, Docquier F, Ozden C. 2009. *Diasporas*. Policy Res. Work. Pap. 4984, World Bank, Washington, DC
- Bergstrand JH. 1985. The gravity equation in international trade: some microeconomic foundations and empirical evidence. *Rev. Econ. Stat.* 67(3):474–81
- Bergstrand JH. 1989. The generalized gravity equation, monopolistic competition, and the factor-proportions theory in international trade. *Rev. Econ. Stat.* 71(1):143–53
- Bergstrand JH, Egger P. 2007. A knowledge- and physical-capital model of international trade flows, foreign direct investment and multinational enterprises. *J. Int. Econ.* 73:278–308
- Bergstrand JH, Egger P. 2011. Gravity equations and economic frictions in the world economy. In *Palgrave Handbook of International Trade*, ed. D Bernhofen, R Falvey, D Greenaway, U Kriekemeier. New York: Palgrave-Macmillan. In press
- Chaney T. 2008. Distorted gravity: the intensive and extensive margins of international trade. *Am. Econ. Rev.* 98:1707–21
- Costinot A, Donaldson D, Komunjer I. 2010. *What goods do countries trade? A quantitative exploration of Ricardo's ideas*. Work. Pap., MIT
- Deardorff A. 1980. The general validity of the law of comparative advantage. *J. Polit. Econ.* 88:941–57
- Eaton B, Kortum S. 2002. Technology, geography, and trade. *Econometrica* 70:1741–79
- Feenstra R. 2004. *Advanced International Trade: Theory and Evidence*. Princeton, NJ: Princeton Univ. Press
- Grogger J, Hanson GH. 2008. *Income maximization and the selection and sorting of international migrants*. NBER Work. Pap. 13821
- Head K, Mayer T, Ries J. 2009. How remote is the off-shoring threat? *Eur. Econ. Rev.* 53:429–44
- Head K, Ries J. 2008. FDI as an outcome of the market for corporate control: theory and evidence. *J. Int. Econ.* 74:2–20
- Helpman E, Melitz MJ, Rubinstein Y. 2008. Estimating trade flows: trading partners and trading volumes. *Q. J. Econ.* 123:441–87
- Helpman E, Melitz MJ, Yeaple SR. 2004. Export versus FDI with heterogeneous firms. *Am. Econ. Rev.* 94:300–16
- Keller W, Yeaple SR. 2009. *Gravity in the weightless economy*. NBER Work. Pap. 15509

- Kleinert J, Toubal F. 2010. Gravity for FDI. *Rev. Int. Econ.* 18(1):1–13
- Leamer EE, Levinsohn J. 1995. International trade theory: the evidence. In *Handbook of International Economics*, Vol. 3, ed. GM Grossman, K Rogoff, pp. 1339–94. Amsterdam: Elsevier Sci. B.V.
- Leamer EE, Stern RM. 1970. *Quantitative International Economics*. Chicago: Aldine
- Martin W, Pham CS. 2008. *Estimating the gravity equation when zero trade flows are frequent*. MPRA Pap. 9453, Univ. Libr. Munich, Ger.
- Martin P, Rey H. 2004. Financial super-markets: Size matters for asset trade. *J. Int. Econ.* 64:335–61
- Matsuyama K. 2007. Beyond icebergs: towards a theory of biased globalization. *Rev. Econ. Stud.* 74:237–53
- McCallum J. 1995. National borders matter: Canada-U.S. regional trade patterns. *Am. Econ. Rev.* 85(3):615–23
- McFadden D. 1974. Conditional logit analysis of qualitative choice behavior. In *Frontiers of Econometrics*, ed. P Zarembka, pp. 105–42. New York: Academic
- Novy D. 2010. *International trade and monopolistic competition without CES: estimating translog gravity*. Unpublished manuscript, Warwick Univ.
- Portes R, Rey H. 2005. The determinants of cross-border equity flows. *J. Int. Econ.* 65:269–96
- Ravenstein EG. 1889. The laws of migration. *J. R. Stat. Soc.* 52(2):241–305
- Redding S, Venables AJ. 2004. Economic geography and international inequality. *J. Int. Econ.* 62(1):53–82
- Samuelson PA. 1952. The transfer problem and transport costs: the terms of trade when impediments are absent. *Econ. J.* 62:278–304
- Santos Silva J, Tenreyro S. 2006. The log of gravity. *Rev. Econ. Stat.* 88(4):641–58
- Savage IR, Deutsch KW. 1960. A statistical model of the gross analysis of transactions flows. *Econometrica* 28:551–72
- Tinbergen J. 1962. *Shaping the World Economy: Suggestions for an International Economic Policy*. New York: Twentieth-Century Fund